PATH PLANNING ALGORITHMS FOR THE ADAPTIVE SENSOR FLEET

Eric Stoneking, Jeff Hosler
NASA Goddard Space Flight Center

The Adaptive Sensor Fleet (ASF) is a fleet of robotic oceangoing surface vessels being developed by NASA in coordination with NOAA. The mission of ASF is to provide the capability for autonomous cooperative survey and sampling of dynamic oceanographic phenomena such as current systems and algae blooms. Each ASF vessel serves as a platform for a variety of instruments. In its present early stages, the assumed instrument platform is the OASIS surface vessel being developed by NASA at the Wallops Flight Facility, but the ASF architecture is being designed for extensibility to accommodate heterogenous fleet elements. This paper describes the path planning algorithms developed for the acquisition phase of a typical ASF task. Given a polygonal target region to be surveyed, the region is subdivided according to the number of vessels in the fleet. The subdivision algorithm seeks a solution in which all subregions have equal area and minimum mean radius. Once the subregions are defined, a dynamic programming method is used to find a minimum-time path for each vessel from its initial position to its assigned region. This path plan includes the effects of water currents as well as avoidance of known obstacles. A fleet-level planning algorithm then shuffles the individual vessel assignments to find the overall solution which puts all vessels in their assigned regions in the minimum time. This "shuffle" algorithm may be described as a process of elimination on the sorted list of permutations of a cost matrix. All these path planning algorithms are facilitated by discretizing the region of interest onto a hexagonal tiling.

INTRODUCTION

The Adaptive Sensor Fleet is a supervisory control system that is designed to use a collection of heterogeneous robotic platforms to optimally perform observations of dynamic environments driven by high-level goals. The ASF provides for observations through high-level goals, supervisory fleet management of robotic platforms (coordination), analysis of environmental science data to use in the decision making process (collaboration), optimal path planning and replanning, identification of science phenomena, and adaptation to dynamic or unknown environments. The concepts addressed in the design of the ASF lay the foundation for a dynamic "Sensor Web" using stationary, surface moving (water or land), airborne, and spacecraft instruments to generate a dynamic network of sensors to achieve the defined science goals.

The ASF is working with the Ocean Atmosphere Sensor Integration System (OASIS) to prototype fleet maneuvering and observation concepts. The OASIS project objective is to provide low-cost ($20K each) mobile, self-navigating surface platforms for ocean sensors, as an alternative to buoys[OASIS]. Benefits of OASIS focus on the ability to navigate to geospatial positions in order to perform in-situ measurements, and then return home to recalibrate and reuse instruments which may otherwise be discarded. These craft will allow scientists to map ocean phenomena that cannot be observed through remote sensing. In addition, these craft will generate significant cost savings for NOAA and NASA especially in the area of data calibration and validation of remote sensing satellite measurements.

OASIS provides a perfect low-cost domain to demonstrate the capability of ASF to perform fleet observations of targets. OASIS craft have several key capabilities that are ideal for concept demonstration:

- Remote Commanding – The OASIS craft can be supplied commands via a remote communication link. They can receive waypoint destinations, science instrument commands, and house-keeping commands. They do not communicate directly with other OASIS craft, but can communicate directly with the control center.
- Semi-autonomous – The OASIS craft can autonomously move to commanded waypoint destinations, and avoid obstacles and shallows in their path.
- Science Collection – The OASIS craft can house instrumentation for in-situ science collection. The science measurements can be transmitted back to the control center through the communication link.
- Heterogeneous – Each OASIS craft will generally have the same guidance and navigation control, but can house instrumentation of different types.

The ASF will use the OASIS craft to demonstrate varying levels of fleet observation concepts. Each observation concept will require unique path planning concepts with varying degrees of difficulty. The ability for ASF to use in-situ measurements in the decision-making process allows the system to adapt to dynamic or unknown environments to optimally complete science observations. This ability is predicated on quality path-planning algorithms and their ability to optimize and replan paths.

Figure 1: The Path Planning Problem

PROBLEM STATEMENT AND PRELIMINARIES

Consider the situation shown in Figure 1. On the right side of the figure is an ocean feature of scientific interest: an algae bloom, perhaps, or a cold eddy. An ASF customer suspects the feature is out there, but doesn't know its precise location or extent. So she identifies a polygonal region of interest, and requests an ASF fleet to survey it. The ASF vessels, beginning from arbitrary initial positions, are tasked to travel quickly and safely to the region, then survey it in some efficient manner. In transit, there are land features, and perhaps other navigational hazards, which will have to be avoided. There are water currents, which need to be accounted for as they affect travel time. Once at the survey region, the survey task must be divided evenly among the fleet. The algorithms which accomplish these tasks are the subject of this paper.

Subdivision of the survey region and vessel minimum-time path planning are vastly simplified by discretization of the problem space. To this end, we overlay a mesh of hexagonal cells over the map. The mesh may be set by the desired resolution of the scientific measurements to be taken, the agility of the vessels, complexity of the terrain, or other factors. Hexagonal cells are preferred over square cells due to their greater directional isotropy; a square cell has four "cardinal" neighbors and four "diagonal" neighbors, while a hex cell has six of each. We number them clockwise from North as shown in figure 2.

Figure 2: Map Cell Neighbor Numbering Scheme

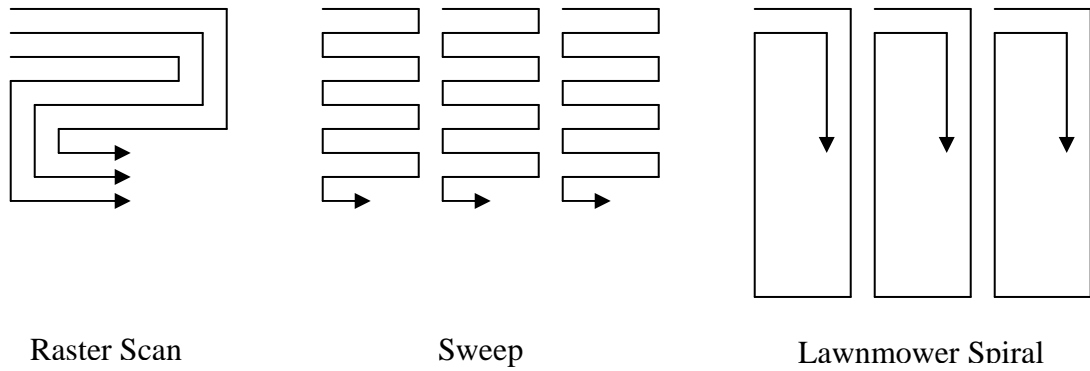Raster Scan   Sweep   Lawnmower Spiral

Figure 3: Some Survey Strategies

All the path planning described in this paper occurs at the initiation of the task. The computation time is negligible when compared to vessel transit times, so adapting to a change in the situation (e.g. an unforeseen obstacle) may be easily accomplished simply by recomputing the path plan as needed, incorporating the new information.

SURVEY REGION SUBDIVISION

Given a region defined by a contiguous convex polygon, we wish to survey it by visiting each cell at least once in as little time as possible. Neglecting water currents and assuming equal speeds for the ASF vessels, it becomes apparent that the minimum-time solution is to divide the region into equal areas, then choose paths which visit each cell with as little retracing as possible. This still leaves a great deal of freedom in selecting a survey strategy. Figure 3 presents three among many candidate schemes: a "raster" scan, a "sweep", and the "lawnmower" spiral. We selected the lawnmower algorithm for development because it rapidly outlines and divides the region of interest, providing a low-resolution survey of the field and then filling in the details. This divide-and-conquer strategy is further refined by choosing the shapes of the subregions to be as "round" as
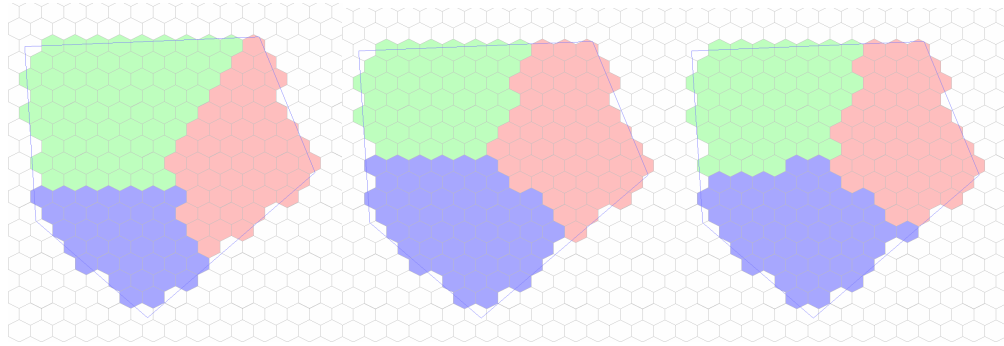
Figure 3: Survey Region Subdivision Sequence

possible. This minimizes the length of the perimeter of a subregion, minimizing the time required to surround and constrict unvisited areas.

The survey region subdivision consists of the following steps:

1. Place N points in region, arranged evenly around the center of the region. These are the initial guesses for the centroids of the subregions.
2. For each cell in the region, find the subregion centroid closest to it. Adjust the measured distances, adding a penalty distance to regions with more member cells (on the first pass, this adjustment is moot, as the subregion memberships haven't been established yet).
3. Assign each cell to the subregion with the shortest adjusted distance.
4. Recompute the centroid of each subregion, and the number of cells in each subregion.
5. Repeat steps 2-4 until the subregion areas are adequately equal.

Figure 3 shows a subdivision sequence. The first iteration chiefly spreads the subregion centroids apart. Subsequent iterations result in exchanges at the boundaries to equalize the subregion areas, with some resultant migration of the centroids. Since the selection criterion is a radius, the subregions tend to roundness, as well as the shape of the survey region allows.

Note that, due to the discretization of the problem, perpetual limit cycles are possible, with a boundary cell being swapped back and forth between two subregions. Suitable convergence criteria are established to guarantee termination of the iteration.

PATH PLANNING

Once the region of interest has been subdivided, the problem becomes one of choosing routes for the fleet vessels which takes them from their (arbitrary) initial positions and places them somewhere within their assigned subregions safely and in some optimal fashion. We first consider the problem of planning a path for a single vessel, and then address the problem of assigning vessels to subregions to achieve a fleet-level optimum solution.
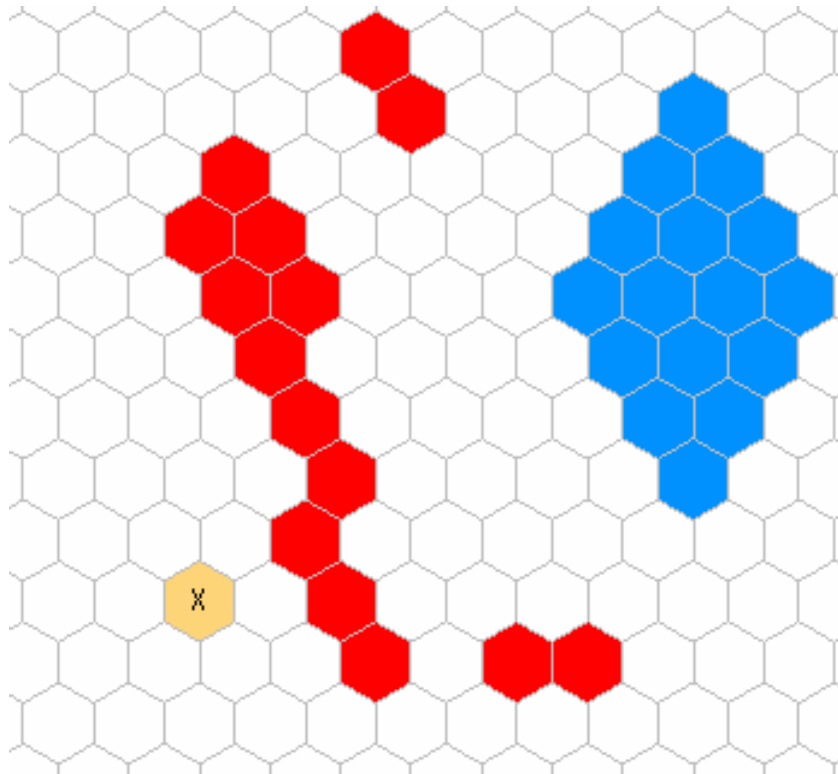
Figure 4: A Simplified Path Planning Problem

Minimum-Time Vessel Path Planning with Obstacles and Water Currents

Consider the situation shown in figure 4. A path, defined by a sequence of waypoints, is to be chosen which connects a vessel's present position (in orange) with a target region (in blue). The path must avoid prohibited areas (e.g. islands, in red), and must minimize a selected cost function (e.g. travel time).

A dynamic programming algorithm [Bryson] is used to find the cost of reaching the target region from every cell on the map:

1. Initialize all map cells as Unresolved. Initialize cells in target region to zero cost. Mark all cells which contain obstacles as Prohibited cells.
2. For each non-prohibited cell in the map,
   a. Find the (12) neighboring cells. Neglecting Prohibited and Unresolved neighbor cells, compute the cost to get to each neighbor cell.
   b. Compute (up to 12) candidate costs by adding the cost to get to each valid neighbor to the cost stored in that neighbor cell.
   c. Select the minimum of these candidate costs, and store this cost in the map cell.
3. Repeat step 2 until a sweep through the entire map changes no cells. This loop is guaranteed to terminate.
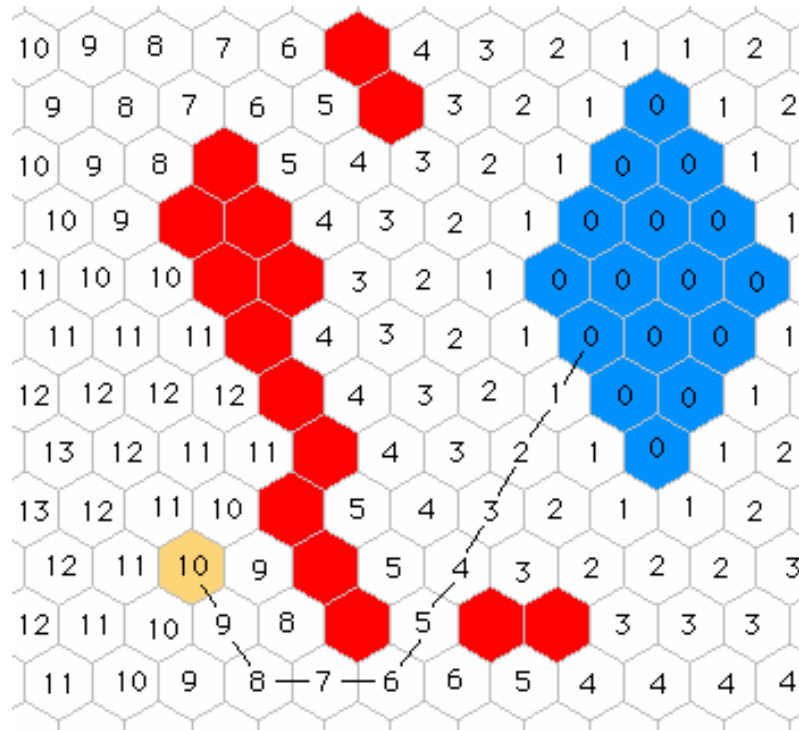
Figure 5: Solution to Simplified Path Planning Problem

Note that prohibited cells are explicitly ruled out. A common alternative is to assign a prohibitively high but finite cost to an obstacle, implicitly forbidding it by pricing it out of the market. This implicit technique is helpful in non-discrete formulations, where analytical cost functions are desirable. A drawback, however, is the risk of not setting the prohibition cost sufficiently high. Faced with a long, circuitous route around an island, such a path planning algorithm may be tempted to take a shortcut over it. Since the present problem is formulated discretely, using explicit prohibition adds no more labor than the implicit technique, and avoids any possibility of "sneaking" through an obstacle rather than going around.

Figure 5 shows a simplified version of the solution. For purposes of illustration, only the six "cardinal" neighbors are considered, and the cost to traverse a cell is simply one unit. With successive iterations, the solved region spreads from the target region until it fills the map.

Once the cost map has been populated, the vessel's path is determined. The first waypoint is set at the vessel's initial position. The next waypoint is placed in the neighboring cell with the lowest time-to-goal cost stored in it. This process is repeated until the target region is reached.
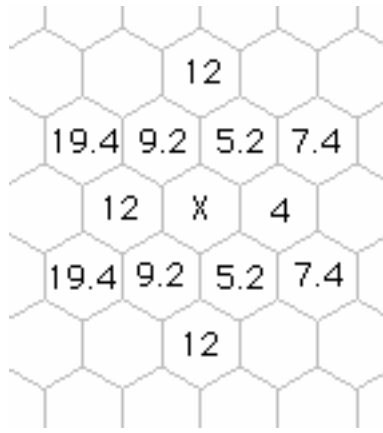
Figure 6: Example of Effect of Water Current on Travel Time

The cost of traveling from a map cell to one of its neighbors is the travel time. When water currents are neglected, the travel time is simply the distance between the cells divided by the speed of the vessel. When water currents are considered, the travel time is affected by the relative direction and speed of the current. Figure 6 illustrates such a situation. In this example, the map cell spacing is 12 km, the vessel's top speed is 2 km/hr, and there is a uniform water current flowing eastward at 1 km/hr. Traveling with the current, the vessel's speed is augmented, and the travel time is lessened. Traveling perpendicular to the current, the vessel must adjust its heading to maintain its desired direction of travel, decreasing its effective speed. And of course, traveling against the current increases the travel time.

Minimum-Time Fleet Path Assignment

Having found a time-optimal solution for a vessel to reach a subregion, we now consider the problem of assigning a vessel to each subregion, such that the time for all vessels to reach their assignments is minimized. There are N! ways to assign N vessels to N subregions. For small values of N (say, N=3), it is cheap to exhaustively check all the permutations, but for only slightly larger fleets, the computational effort to do so becomes unreasonable (3,628,800 permutations for a perfectly reasonable fleet of 10 vessels). We developed a more parsimonious assignment procedure which scales with N^2 rather than with N!.

This procedure is best explained by an artificial example. Three vessels are to be assigned to three survey subregions. The associated costs have been computed and arranged in the matrix C, each element C(i,j) being the cost for vessel *i* to reach subregion *j*.

$$C = \begin{matrix} 6.3 & 4.9 & 7.1 \\ 9.3 & 0.5 & 1.1 \\ 5.3 & 0.4 & 8.1 \end{matrix}$$

We wish to select three elements of C so that 1) no two of the three are in the same row or column, which ensures that no two vessels are assigned to the same region, and no two regions are assigned to the same vessel; and 2) the largest value of the three is as small as

we can make it, which minimizes the time for all vessels to be in position. So we form a list, sorted by decreasing cost, and start a process of elimination:

Rank	Cost	Element
1	9.3	(2,1)
2	8.1	(3,3)
3	7.1	(1,3)
4	6.3	(1,1)
5	5.3	(3,1)
6	4.9	(1,2)
7	1.1	(2,3)
8	0.5	(2,2)
9	0.4	(3,2)

Elimination of the three highest-cost elements results in:

$$C = \begin{array}{ccc} 6.3 & 4.9 & \cancel{7.1} \\ \cancel{9.3} & 0.5 & 1.1 \\ 5.3 & 0.4 & \cancel{8.1} \end{array}$$

Element (2,3) is the last surviving element in its column, so it must be part of the solution. We denote this with bold typeface. Since it is part of the solution, no other element in row 2 can be, so we eliminate element (2,2):

$$C = \begin{array}{ccc} 6.3 & 4.9 & \cancel{7.1} \\ \cancel{9.3} & \cancel{0.5} & \mathbf{1.1} \\ 5.3 & 0.4 & \cancel{8.1} \end{array}$$

This action results in no new Last-in-Column or Last-in-Row conditions, so we return to our sorted list and eliminate the largest surviving member, (1,1):

$$C = \begin{array}{ccc} \cancel{6.3} & 4.9 & \cancel{7.1} \\ \cancel{9.3} & \cancel{0.5} & \mathbf{1.1} \\ 5.3 & 0.4 & \cancel{8.1} \end{array}$$

Eliminating element (1,1) leaves (1,2) as the last in its row, and (3,1) as last in its column. They must be part of the solution, and element (3,2), since it shares a row (and a column) with a solution member, must not be in the solution:

$$C = \begin{array}{ccc} \cancel{6.3} & \mathbf{4.9} & \cancel{7.1} \\ \cancel{9.3} & \cancel{0.5} & \mathbf{1.1} \\ \mathbf{5.3} & \cancel{0.4} & \cancel{8.1} \end{array}$$

This algorithm may be summarized:
1. Compute the NxN cost matrix, C.
2. Find the list of elements, sorted in order of decreasing cost.
3. While any element in C is neither eliminated nor selected as part of the solution:
    a. While any element is Last-in-Row (LiR) or Last-in-Column (LiC)
        i. Select LiR or LiC element as part of solution.
        ii. Eliminate remainder of row and column.
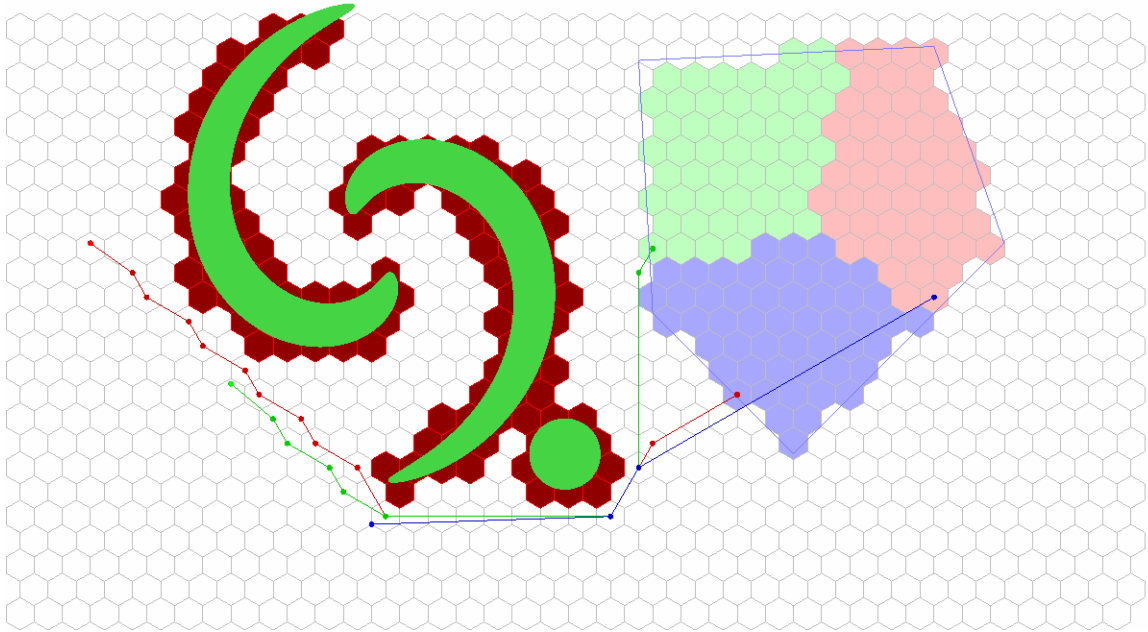   b. Eliminate highest-cost element in sorted list.

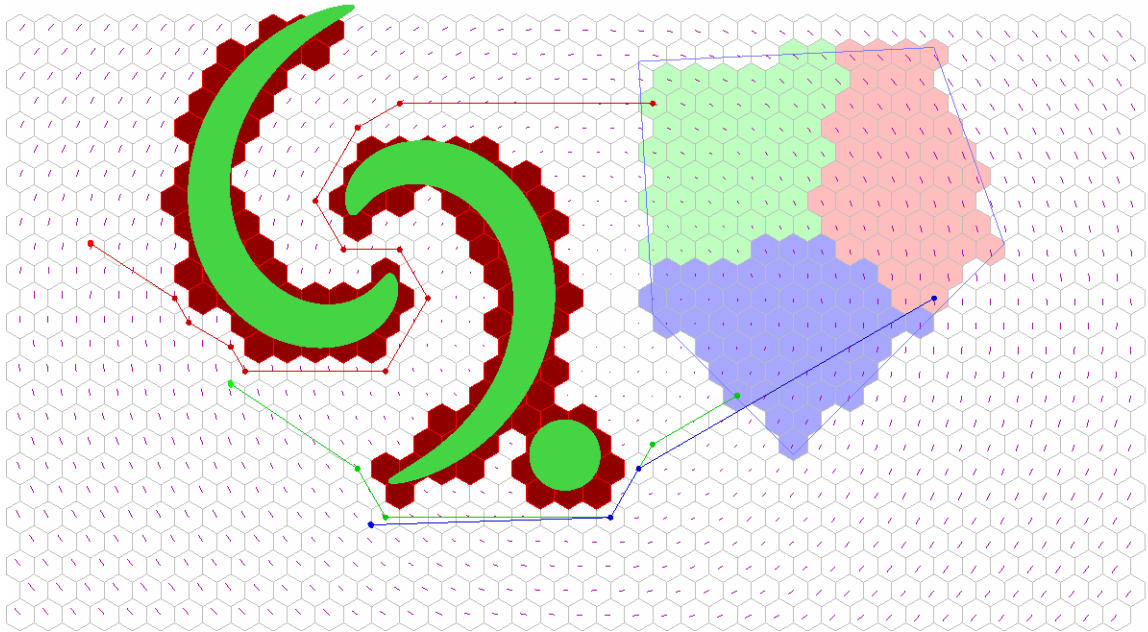**Figure 7: Vessel Paths With No Water Currents**

Figure 8: Vessel Paths With Water Currents

Note that the first N-1 elements on the sorted list may be eliminated directly, since they can never be the Last-in-Row or Last-in-Column.

Figures 7 and 8 show the paths assigned for our case study. Figure 7 shows the paths assigned neglecting water currents, and figure 8 shows the solution with a water current field applied. Introduction of water current knowledge makes the inter-island

channel a more economical path for one vessel, which in turn affects the vessels' assigned subregions.

CONCLUSION

In this paper we have presented three of the path planning algorithms used on the Adaptive Sensor Fleet project at NASA Goddard Space Flight Center. In combination, these algorithms enable a fleet of water surface vessels to safely and efficiently reach a defined region from arbitrary initial positions in preparation for a cooperative survey of the region. Discretization of the problem simplifies map management and allows straightforward and robust obstacle avoidance. The map is tiled with hexagonal cells to preserve as much directional isotropy as possible.

Division of the survey area is accomplished with an iterative area-exchange method. Path planning for a fleet member uses dynamic programming to reach the target area in minimum time, in the presence of obstacles and water currents. Vessel assignments are shuffled to yield a fleet path plan which places all vessels in their assigned subregions in minimum time.

The ASF project is currently operating in a simulated environment. OASIS craft simulators are being used for guidance, navigation and control, and the Regional Ocean Modeling System (ROMS) is being used to provide the ASF with modeled ocean measurements. Water tests with the physical OASIS craft are planned for July 2005. These tests will demonstrate many of the concepts described in this paper, including fleet coordination and supervisory-level collaboration.

The goal of the ASF is to implement a "Sensor Web" architecture combining heterogeneous elements: stationary platforms, water and land surface vessels, airborne and spaceborne instruments. Although the immediate application is to study the Earth's oceans, the architecture and techniques developed and demonstrated by ASF will be readily applicable to exploration of other planets by landers, rovers, and orbiting platforms.

REFERENCES
[OASIS] Troy Ames, Dr. John Moisan, Pam Pittman. "Ocean-Atmosphere Sensor Integration System (OASIS) Prototype Report", NASA Goddard Space Flight Center, 2003.
[BRYSON] Arthur E. Bryson, Jr. and Yu-Chi Ho. *Applied Optimal Control*. Hemisphere Publishing Company, New York, 1975.